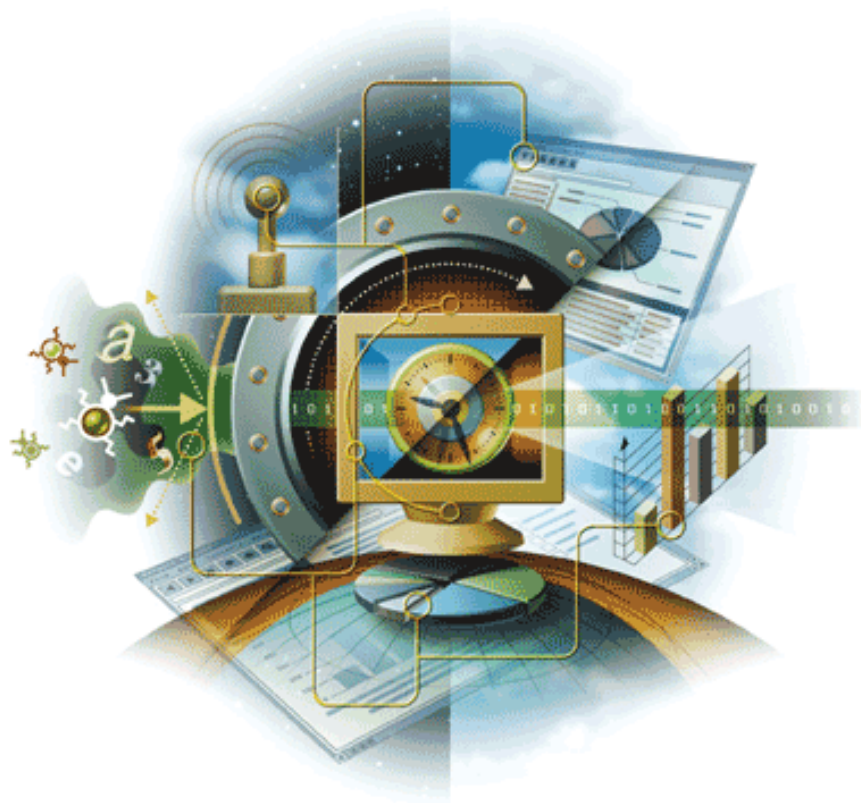


McAfee® IntruShield® IPS System

Best Practices



McAfee®
Network Protection

Industry-leading intrusion prevention solutions

McAfee®

COPYRIGHT

© 2002 - 2005 McAfee, Inc. All Rights Reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form or by any means without the written permission of McAfee, Inc., or its suppliers or affiliate companies. To obtain this permission, write to the attention of the McAfee, Inc. legal department at: 5000 Headquarters Drive, Plano, Texas 75024, or call +1-972-963-8000.

TRADEMARK ATTRIBUTIONS

Active Firewall, Active Security, ActiveSecurity (in Katakana), ActiveHelp, ActiveShield, AntiVirus Anyware and design, Bomb Shelter, Certified Network Expert, Clean-Up, CleanUp Wizard, ClickNet, CNX, CNX Certification Certified Network Expert and design, Covert, Design (Stylized E), Design (Stylized N), Disk Minder, Distributed Sniffer System, Distributed Sniffer System (in Katakana), Dr Solomon's, Dr Solomon's label, Entercept, Enterprise SecureCast, Enterprise SecureCast (in Katakana), ePolicy Orchestrator, EZ SetUp, First Aid, ForceField, GMT, GroupShield, GroupShield (in Katakana), Guard Dog, HomeGuard, Hunter, IntruShield, Intrusion Prevention Through Innovation, IntruVert Networks, LANGuru, LANGuru (in Katakana), M and Design, McAfee, McAfee (in Katakana), McAfee and design, McAfee.com, McAfee VirusScan, NA Network Associates, Net Tools, Net Tools (in Katakana), NetCrypto, NetOctopus, NetScan, NetShield, NetStalker, Network Associates, Network Associates Coliseum, NetXray, NotesGuard, Nuts & Bolts, Oil Change, PC Medic, PCNotary, PrimeSupport, Recoverkey, Recoverkey - International, Registry Wizard, RingFence, Router PM, SecureCast, SecureSelect, Sniffer, Sniffer (in Hangul), SpamKiller, Stalker, TIS, TMEG, Total Network Security, Total Network Visibility, Total Network Visibility (in Katakana), Total Virus Defense, Trusted Mail, UnInstaller, Virex, Virus Forum, ViruScan, VirusScan, WebScan, WebShield, WebShield (in Katakana), WebSniffer, WebStalker, WebWall, What's The State Of Your IDS?, Who's Watching Your Network, WinGauge, Your E-Business Defender, Zip Manager are registered trademarks or trademarks of McAfee, Inc. and/or its affiliates in the US and/or other countries. Sniffer® brand products are made only by McAfee, Inc. All other registered and unregistered trademarks herein are the sole property of their respective owners.

PATENTS

Protected by US Patents 5,361,359; 5,557,742; 6,275,942; 6,301,699; 6,412,071; 6,496,875; 6,510,448; 6,513,122; 6,546,493; 6,584,508; 6,587,888; 6,668,289; 6,684,329

LICENSE INFORMATION

License Agreement

NOTICE TO ALL USERS: CAREFULLY READ THE APPROPRIATE LEGAL AGREEMENT CORRESPONDING TO THE LICENSE YOU PURCHASED, WHICH SETS FORTH THE GENERAL TERMS AND CONDITIONS FOR THE USE OF THE LICENSED SOFTWARE. IF YOU DO NOT KNOW WHICH TYPE OF LICENSE YOU HAVE ACQUIRED, PLEASE CONSULT THE SALES AND OTHER RELATED LICENSE GRANT OR PURCHASE ORDER DOCUMENTS THAT ACCOMPANIES YOUR SOFTWARE PACKAGING OR THAT YOU HAVE RECEIVED SEPARATELY AS PART OF THE PURCHASE (AS A BOOKLET, A FILE ON THE PRODUCT CD, OR A FILE AVAILABLE ON THE WEB SITE FROM WHICH YOU DOWNLOADED THE SOFTWARE PACKAGE). IF YOU DO NOT AGREE TO ALL OF THE TERMS SET FORTH IN THE AGREEMENT, DO NOT INSTALL THE SOFTWARE. IF APPLICABLE, YOU MAY RETURN THE PRODUCT TO MCAFFEE, INC. OR THE PLACE OF PURCHASE FOR A FULL REFUND.

Attributions

This product includes or may include:

◆ Some software programs that are licensed (or sublicensed) to the user under the GNU General Public License (GPL) or other similar Free Software licenses which, among other rights, permit the user to copy, modify and redistribute certain programs, or portions thereof, and have access to the source code. The GPL requires that for any software covered under the GPL which is distributed to someone in an executable binary format, that the source code also be made available to those users. For any such software covered under the GPL, the source code is made available on this CD. If any Free Software licenses require that McAfee, Inc. provide rights to use, copy or modify a software program that are broader than the rights granted in this agreement, then such rights shall take precedence over the rights and restrictions herein.
◆ Software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).
◆ Cryptographic software written by Eric A. Young (eaay@cryptsoft.com) and software written by Tim J. Hudson (tjh@cryptsoft.com).
◆ OpenSSH software copyrighted by University of California, © 1980, 1993 The Regents of the University of California.
◆ Software copyrighted by Markus Friedl, Theo de Raadt, Niels Provos, Dug Song, Aaron Campbell, Damien Miller, Kevin Steves.
◆ Software copyrighted CORE SDI S.A., © 1998 Buenos Aires, Argentina.
◆ Software copyrighted by David Mazieres, © 1995, 1996.
◆ Software copyrighted by Gary S. Brown, © 1986.
◆ Software copyrighted by Tatu Ylonen, © 1995.
◆ Net-SNMP software copyrighted by Carnegie Mellon University; © 1989, 1991, 1992, and The Regents of the University of California © 1996, 1998-2000.
◆ Software copyrighted by Cambridge Broadband Ltd., © 2001, All rights reserved.
◆ MySQL software copyrighted by MySQL, AB.
◆ Software developed and copyrighted by the Apache Software Foundation, © 1999, 2000, which software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation and was originally based on software copyright © 1999, International Business Machines, Inc., < http://www.ibm.com >.
◆ JOnAS software originally created and copyrighted by Bull S.A., © 1999 Bull S.A.
◆ SUN Java Secure Socket Extension software copyrighted by Software copyrighted by Sun Microsystems®, Inc.
◆ Software copyrighted by RSA Data Security, Inc.
◆ Software copyrighted by Lumos Technologies, © 2000-2002 Lumos Technologies, Inc. All rights reserved.
◆ Software owned by Broadcom Corporation which is licensed by not sold. Title to and ownership of these portions and any portion thereof remain with Broadcom Corporation. All express and implied warranties are disclaimed on behalf of Broadcom. Broadcom and its licensors disclaim any liability for any special, indirect, exemplary, incidental or consequential damages.
◆ Software copyrighted by Wind River Systems, Inc., © 2000.
◆ Software developed by Cypress Semiconductor Corporation. All rights to the Cypress Semiconductor Corporation development software belong to Cypress Semiconductor Corporation, owner of the former Lara Technology, Inc. No part of this Cypress Semiconductor Corporation development software code may be used, modified, reproduced, transmitted, distributed, displayed or translated without the express written consent of Cypress Semiconductor Corporation. Proprietary legends within such code are to be preserved in all copies thereof. Cypress Semiconductor Corporation makes no warranty of any kind, express or implied, with regards to this code, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.
◆ Software copyrighted by PMC-Sierra, Inc., © 1998-2002. All rights to the PMC-Sierra development software belongs to PMC-Sierra, Inc. No part of this code may be used, modified, reproduced, transmitted, distributed, displayed or translated without the express written consent of PMC-Sierra, Inc.
◆ Software copyrighted by Vitesse Semiconductor Corporation, © 1997-2000. No part of this Vitesse Semiconductor Corporation code may be used, modified, reproduced, transmitted, distributed, displayed or translated without the express written consent of Vitesse Semiconductor Corporation.
◆ Software copyrighted by Todd C. Miller, (c) 1998.
◆ Software copyrighted by The Regents of the University of California, (c) 1990, 1993, with code derived from software contributed to Berkeley by Chris Torek.

Table of Contents

1	IntruShield Best Practices	1
	Contents of this document	1
	Pre-installation considerations	2
	Hardware requirements	2
	Determining your database size	2
	Deploying a large number of sensors	3
	Considerations for large deployments	3
	Staging sensors prior to deployment	4
	Deploying sensors in phases	4
	Facilitating troubleshooting	4
	Ensuring connectivity between the sensor and other network devices	5
	Duplex mismatches	5
	Valid auto-negotiation and speed configurations	6
	Gigabit auto-negotiation (no link to connected device)	6
	Troubleshooting a Duplex Mismatch with Cisco Devices	6
	Cisco PIX® Firewall	7
	Cisco CSS 11000	7
	Cisco Catalyst® 2900XL, 3500XL Series (Hybrid)	7
	Cisco Catalyst 4000, 5000, 6000 Series (Native)	7
	Cisco IOS® for Catalyst 4000, 6000 Series	7
	Explanation of CatOS show port command counters	7
	Auto-negotiation	9
	Situations that may lead to auto-negotiation Issues	9
	Initial tuning	9
	High-volume attacks	9
	Alert filters	10
	DoS	10
	Sensor actions	11
	Response management	11
	Creating rule sets	12
	Default Inline IPS	12
	Port clustering on asymmetric networks	12
	Maintenance, backup, and database tuning	13
	Alerts and disk space maintenance	14
	Alert states	14
	Dbtuning.bat	14
	Purge.bat	15
	Using the File Maintenance scheduler	15
	Backup	15
	Access Control Lists (ACL)	16
	Working on performance issues	17
	Sniffer trace	17
	Understanding data link errors	17
	Half-duplex setting	17
	Full-duplex setting	17
	SSL best practices	18
	SSL only traffic - throughput	18
	SSL only traffic - throughput	18
	SSL traffic mixed with HTTP 1.1 traffic	18
	Supported Web servers	19

Unsupported SSL functionality. 20

Sensor performance with HTTP Response processing. 20

 HTTP Response processing enabled for both inbound and outbound traffic . 20

 HTTP Response processing enabled in one direction only (inbound
 OR outbound) 21

1

IntruShield Best Practices

This document discusses recommended practices for using IntruShield most effectively. Topics covered include installation, tuning, rule set creation, connectivity and maintenance. For in depth information about IntruShield IPS protection, see the following documents:

- *IntruShield Manager Installation Guide*
- *IntruShield Getting Started Guide*
- *IntruShield Manager Administrator's Guide*
- *IntruShield Troubleshooting Guide*
- *IntruShield Release Notes*

Contents of this document

- *Pre-installation considerations on page 2*
- *Deploying a large number of sensors on page 3*
- *Facilitating troubleshooting on page 4*
- *Ensuring connectivity between the sensor and other network devices on page 5*
- *Initial tuning on page 9*
- *Creating rule sets on page 12*
- *Port clustering on asymmetric networks on page 12*
- *Maintenance, backup, and database tuning on page 13*
- *Access Control Lists (ACL) on page 16*
- *Working on performance issues on page 17*
- *SSL best practices on page 18*
- *Sensor performance with HTTP Response processing on page 20*

Pre-installation considerations

Hours and even days can be saved during the IntruShield installation and tuning process if you are fully prepared. The *IntruShield Troubleshooting Guide* spells out the list of tasks that you should complete before you schedule your IntruShield Manager software installation.



The *IntruShield Troubleshooting Guide* is a new document as of release 3.1; however, most of the techniques described in the document apply to all versions of IntruShield.

Hardware requirements

The larger your deployment, the more high-end your Manager server should be. Many IntruShield issues result from an underpowered Manager server.

For example, to manage 40 or more sensors, we recommend going beyond the hardware recommended in the release notes. The following is a recommended minimum hardware configuration:

Manager server with embedded MySQL database

- 4GB RAM
- 2 x 3.2 Pentium processors
- 80 GB hard disk space (or greater)



You will experience better performance in your configuration and data forensic tasks by connecting to the Manager from a browser on a client machine. Performance may be slow if you connect to the Manager using a browser on the server machine itself.

Determining your database size

The amount of space your database will require is governed by many factors, most of which are unique per deployment. They boil down to how much data you want to retain in the database and for how long.

Things to consider when planning are:

- **Aggregate alert and packet log volume from all sensors**—Many sensors equals higher alert volume and will require additional storage capacity. Note that an alert is roughly 200 bytes on average, while a packet lot is approximately 450 bytes.
- **Lifetime of alert and packet log data:** How long before you archive and then delete an alert? Maintaining your data for a long period of time (e.g., one year) will require additional storage capacity to accommodate both old and new data.

As a best practice, McAfee recommends archiving and deleting old alert data regularly and attempting to keep your active database to about 40GB.



You can find capacity planning information in the *Manager Administrator's Guide*, Appendix B.

Deploying a large number of sensors

What is a “large number of sensors?” For the purpose of this document, we’ll break down deployment size into Small, Medium, Large, and Very Large.

Size	Number of Sensors
Small	Fewer than 10
Medium	10 to 35
Large	36 to 70
Very Large	more than 71

Considerations for large deployments

You will need to develop your own practices for tasks such as sensor deployment, sensor software upgrades, and signature set updates. McAfee recommends that you consider these tasks up front and establish written standard operating procedures on how you prefer to accomplish the task in your environment *before* you deploy.

- **Signature set downloads** - Both signature set downloads are applied to sensors serially, not concurrently. For releases prior to 3.1, where many performance improvements were made to the signature set download process, the process would take approximately 3 minutes per sensor. As a user, you must decide the point at which a signature set update process becomes too time consuming, at which we would recommend utilizing a second ISM. For example, in a deployment with 50 sensors, with versions prior to 3.1, it will take approximately 2.5 hours to complete a signature set update. The same is true for policy updates. (Again, with version 3.1, this process has been reduced to a matter of minutes, not hours.)
- **Sensor software updates** - While signature set updates very rarely require a sensor reboot, all sensor software updates do require a reboot. A reboot can take up to 5 minutes. You can schedule this process, but any update from the Manager causes the process to take place sequentially, one sensor at-a-time. You can instead use the TFTP method for updating the sensor image, enabling you to load an image on the sensor via the sensor’s CLI, and thus load images concurrently (and thus faster).

The process of using TFTP to update your sensor software is documented in the *Sensor Configuration Guide*.

- **Usability** - Depending on the number of VIPS and Admin Domains utilized in your deployment, the ISM Resource Tree can become very crowded, which can lead to more scrolling within the view into your deployment. It can also lead to confusion if you have not provided unique, recognizable names for your sensors and any VIDS you create. The name appears both in the Resource Tree of the Manager, but also in alert data and reports, for example, and providing vague names can lead to quite a bit of confusion. For example, compare a worldwide deployment where sensors are named “4010-1” through “4010-25” as opposed to “UK-London-sens1,” “India-Bangalore-sens1,” and so on. Your VIDS names should also be clear and easy for everyone maintaining the network to recognize at a glance.

- **Alert Traffic** - if “chatty” policies are deployed on the sensors, there is potential to starve ISM resources as the resulting alerts are passed to the Manager. The more sensors with high-volume alerting, the more data you will have to sift through as you tune your policies.
- **Start up load on the Manager** - When the ISM starts, establishing connections with all sensors can be time consuming, as sensors continue to collect alerts while communication with the Manager is lost, and each sensor must then pass its alert data to the Manager when connectivity is re-established.
- **Concurrent processes** - Be aware of the time periods in which your scheduled processes (such as database backup or report generation) occur, and try not to attempt other tasks during that time period, as this can lead to process locking. This includes having many users logged into the system simultaneously.

Staging sensors prior to deployment

With large or very large deployments, and/or if you are planning to release sensors to various geographical regions or difficult to reach locations, you may want to consider staging your sensors before you release them to their final destination. For example, use a Manager in a lab environment to push sensor software to the sensor, bring up the sensor to establish that it is working to your satisfaction, and then box the configured sensor and send it to its final destination. Or you might use the TFTP feature to load the sensor image at one location before shipping the sensor to another.

Deploying sensors in phases

Most IntruShield customers begin their deployment in their lab environment; here they test the sensor functionality, familiarize themselves with the Manager, create an initial policy, and once they are comfortable with the product, then they deploy the sensor into a live environment. The first sensor is always the slowest one to be deployed. McAfee provides a few recommendations for this process:

- Spend time creating effective policies before you deploy. Having more data available makes the tuning process easier, but policies like the IntruShield provided All-Inclusive policy can overwhelm you with data if every sensor in a large deployment is running it without any customization.
- Stagger your sensor deployment in phases. As each new batch of sensors provides you with more data points, you can tune your policies more effectively and become more aggressive in the number of sensors you deploy in the next phase.

Facilitating troubleshooting

When an in-line device experiences problems, one instinct is to physically pull it out of the path; to disconnect the cables and let traffic flow unimpeded while the device can be examined elsewhere. McAfee recommends you first try the following techniques to troubleshoot a sensor issue:

- All sensors have a *Layer2 Passthru* feature. If you feel your sensor is causing network disruption, before you remove it from the network, issue the following command:

```
layer2 mode assert
```


This pushes the sensor into Layer2 Passthru (L2) mode, causing traffic to flow through the sensor while bypassing the detection engine. Check to see whether your services are still affected; if they are, then you have eliminated certain sensor hardware issues; the problem could instead be a network issue or a configuration issue. (The `layer2 mode deassert` command pushes the sensor back to detection mode.)

- McAfee recommends that you configure Layer2 Passthru Mode on each sensor. This enables you to set a threshold on the sensor that pushes the sensor into L2 bypass mode if the sensor experiences a specified number of errors within a specified timeframe. Traffic then continues to flow directly through the sensor without passing to the detection engine.
- Connect a fail-open kit, which consists of a bypass switch and a controller, to any GE monitoring port pairs on the sensor. If a kit is attached to the sensor, disabling the sensor ports forces traffic to flow through the bypass switch, effectively pulling the sensor out of the path. For FE monitoring ports, there is no need for the external kit. Sensors with FE ports contain an internal tap; disabling the ports will send traffic through the internal tap, providing fail-open functionality.



Note that the sensor will need to reboot to move out of L2 mode only if the sensor entered L2 mode because of internal errors. (It does not need a reboot if the `layer2 mode assert` command was used to put the sensor into L2 mode).

A sensor reboot breaks the link connecting the devices on either side of the sensor and requires the renegotiation of the network link between the two devices surrounding the sensor.

Depending on the network equipment, this disruption should range from a couple of seconds to more than a minute with certain vendors' devices.

A very brief link disruption might occur while the links are renegotiated to place the sensor back in in-line mode.

Ensuring connectivity between the sensor and other network devices

The most common sensor deployment problems relate to configuration of the monitoring port speed and duplex settings. Speed determination issues may result in no connectivity between the sensor and its network device partners on either side.

Duplex mismatches

A duplex mismatch (e.g., one end of the link in full-duplex and the other in half-duplex) may result in performance issues, intermittent connectivity, and loss of communication. It can also create subtle problems in applications. For example, if a Web server is talking to a database server through an Ethernet switch with a duplex mismatch, small database queries may succeed, while large ones fail due to a timeout.

Manually setting the speed and duplex to full-duplex on only one link partner generally results in a mismatch. This common issue results from disabling auto-negotiation on one link partner and having the other link partner default to a half-duplex configuration, creating the mismatch. This is the reason why speed and duplex cannot be hard-coded on only one link partner. If your intent is not to use auto-negotiation, you must manually set both link partners' speed and duplex settings to full-duplex.

Valid auto-negotiation and speed configurations

The table below summarizes all possible settings of speed and duplex for IntruShield sensors and switch ports.

IntruShield Configuration 10/100 port (Speed/Duplex)	Configuration of Switch (Speed/Duplex)	Resulting Sensor Speed/Duplex	Resulting Catalyst Speed/Duplex	Comments
100 Mbps Full-duplex	1000 Mbps Full-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch
100 Mbps Full-duplex	AUTO	100 Mbps Full-duplex	100 Mbps Full-duplex	Duplex Mismatch 1
100 Mbps Full-duplex	1000 Mbps Full-duplex	100 Mbps Full-duplex	100 Mbps Full-duplex	Correct Manual Configuration2
100 Mbps Half-duplex	AUTO	100 Mbps Half-duplex	100 Mbps Half-duplex	Link is established, but switch does not see any auto-negotiation information from IntruShield and defaults to half-duplex when operating at 10/100 Mbps.
10 Mbps Half-duplex	AUTO	100 Mbps Half-duplex	100 Mbps Half-duplex	Link is established, but switch does not see Fast Link Pulse (FLP) and defaults to 10 Mbps half-duplex.
10 Mbps Half-duplex	1000 Mbps Half-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch.

Gigabit auto-negotiation (no link to connected device)

Gigabit Ethernet has an auto-negotiation procedure that is more extensive than that which is used for 10/100 Mbps Ethernet (per Gigabit auto-negotiation specification IEEE 802.3z-1998). The Gigabit auto-negotiation negotiates flow control, duplex mode, and remote fault information. You must either enable or disable link negotiation on both ends of the link. Both ends of the link must be set to the same value or the link will not connect.

If either device does not support Gigabit auto-negotiation, disabling Gigabit auto-negotiation forces the link up.

Troubleshooting a Duplex Mismatch with Cisco Devices

When troubleshooting connectivity issues with Cisco switches or routers, verify that the sensor and the switch/routers are using a valid configuration. The `show intfport <port>` command on the IntruShield sensor CLI will help reveal errors.

Sometimes there are duplex inconsistencies between IntruShield and the switch port. Symptoms include poor port performance and frame check sequence (FCS) errors that increment on the switch port. To troubleshoot this issue, manually configure the switchport to 100 Mbps, half-duplex. If this action resolves the connectivity problems, you may be running into this issue. Contact Cisco's TAC for assistance.

Use the following commands to verify fixed interface settings on some Cisco devices that connect to IntruShield sensors:

Cisco PIX® Firewall

- interface ethernet0 100full

Cisco CSS 11000

- interface ethernet-3
 - phy 100Mbps-FD

Cisco Catalyst® 2900XL, 3500XL Series (Hybrid)

- interface FastEthernet0/2
 - duplex full
 - speed 100

Cisco Catalyst 4000, 5000, 6000 Series (Native)

- set port speed 1/1 100
- set port duplex 1/1 full

Cisco IOS® for Catalyst 4000, 6000 Series

- Router(config)# interface fastethernet slot/port
- Router(config-if)# speed 100
- Router(config-if)# duplex full

When troubleshooting IntruShield performance issues with Cisco switches, view the output of the `show port mod/port` command, and note the counter information.

Explanation of CatOS show port command counters

Counter	Description	Possible Causes
Alignment Errors	Alignment errors are a count of the number of frames received that do not end with an even number of octets and have a bad CRC.	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device generating frames that do not end with on an octet and have a bad FCS.
FCS	FCS error count is the number of frames that were transmitted or received with a bad checksum (CRC value) in the Ethernet frame. These frames are dropped and not propagated onto other ports.	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device generating frames with bad FCS.
Xmit-Err	This is an indication that the internal transmit buffer is full.	This is an indication of excessive input rates of traffic. This is also an indication of transmit buffer being full. The counter should only increment in situations in which the switch is unable to forward out the port at a desired rate. Situations such as excessive collisions and 10 Mb ports cause the transmit buffer to become full. Increasing speed and moving the link partner to full-duplex should minimize this occurrence.

Counter	Description	Possible Causes
Rcv-Err	This is an indication that the receive buffer is full.	This is an indication of excessive output rates of traffic. This is also an indication of the receive buffer being full. This counter should be zero unless there is excessive traffic through the switch. In some switches, the Out-Lost counter has a direct correlation to the Rcv-Err.
UnderSize	These are frames that are smaller than 64 bytes (including FCS) and have a good FCS value.	This is an indication of a bad frame generated by the connected device.
Single Collisions	Single collisions are the number of times the transmitting port had one collision before successfully transmitting the frame to the media.	This is an indication of a half-duplex configuration.
Multiple Collisions	Multiple collisions are the number of times the transmitting port had more than one collision before successfully transmitting the frame to the media.	This is an indication of a half-duplex configuration.
Late Collisions	A late collision occurs when two devices transmit at the same time and neither side of the connection detects a collision. The reason for this occurrence is that the time to propagate the signal from one end of the network to another is longer than the time to put the entire packet on the network. The two devices that cause the late collision never see that the other is sending until after it puts the entire packet on the network. Late collisions are detected by the transmitter after the first time slot of the 64-byte transmit time occurs. They are only detected during transmissions of packets longer than 64 bytes. Its detection is exactly the same as it is for a normal collision; it just happens later than it does for a normal collision.	This is an indication of faulty hardware (NIC, cable, or switch port) or a duplex mismatch.
Excessive Collisions	Excessive collisions are the number of frames that are dropped after 16 attempts to send the packet resulted in 16 collisions.	This is an indication of overutilization of the switch port at half-duplex or duplex mismatch.
Carrier Sense	Carrier sense occurs every time an Ethernet controller wants to send data and the counter is incremented when there is an error in the process.	This is an indication of faulty hardware (NIC, cable, or switch port).
Runts	These are frames smaller than 64 bytes with a bad FCS value.	This is an indication of the result of collisions, duplex mismatch, IEEE 802.1Q (dot1q), or an Inter-Switch Link Protocol (ISL) configuration issue.
Giants	These are frames that are greater than 1518 bytes and have a bad FCS value.	This is an indication of faulty hardware, dot1q, or an ISL configuration issue.

Auto-negotiation

Auto-negotiation issues typically do not result in link establishment issues. Instead, auto-negotiation issues mainly result in a loss of performance. When auto-negotiation leaves one end of the link in, for example, full-duplex mode and the other in half-duplex (also known as a duplex mismatch), errors and retransmissions can cause unpredictable behavior in the network. This can cause performance issues, intermittent connectivity, and loss of communication. Generally these errors are not fatal—traffic still makes it through—but locating and fixing them is a time-waster.

Situations that may lead to auto-negotiation Issues

Auto-negotiation issues with the IntruShield sensor may result from nonconforming implementation, hardware incapability, or software defects.

Generally, if the switch used with the sensor adheres to IEEE 802.3u auto-negotiation specifications and all additional features are disabled, auto-negotiation should properly negotiate speed and duplex, and no operational issues should exist.

- Problems may arise when vendor switches/routers do not conform exactly to the IEEE specification 802.3u.
- Vendor-specific advanced features that are not described in IEEE 802.3u for 10/100 Mbps auto-negotiation (such as auto-polarity or cabling integrity) can also lead to hardware incompatibility and other issues.

Initial tuning

As of software version 2.1, all sensors, on initial deployment, have the Intrushield 'Default Inline IPS' policy loaded on all interfaces. The "Default Inline IPS" policy can be changed at the Root admin layer. McAfee recommends, where appropriate, to use this or another IntruShield-provide policy as a starting point, but to tune these into segment-tailored custom policies. These tailored policies can be either cloned versions of Intrushield pre-configured policies or custom-built policies that employ custom rule sets. An appropriately tuned policy will reduce false positives.

Though each network environment has unique characteristics, the following best practices can make tuning more efficient and effective.

High-volume attacks



Note that as you interact with Intrushield policies, you encounter the term "attack," not "signature." Intrushield defines an attack as being comprised of one or more signatures, thresholds, anomaly profiles, or correlation rules, where each method is used to detect an attempt to exploit a particular vulnerability in a system. These signatures and checks may contain very specific means for identifying a specific known exploit of the vulnerability, or more generic detection methods that aid in detecting unknown exploits for the vulnerability.

Take attacks that are generating the most alerts (use the **top10** table in the **Consolidated View** within **Alert Viewer**) and investigate their legitimacy. See Chapter 11 of the *Manager Administrator's Guide* for more details.

Many of the top alerts seen on the initial deployment of a sensor will be common false positives seen in many environments. Typically, at the beginning of the tuning process, it will be evident that your network or security policy will affect the overall level of alerts. If, for instance, AOL IM is allowed traffic on the network then there might not be a need to alert on AOL IM set-up flows.

Alert filters

When a particular alert is declared a false positive, the next decision is whether to disable the corresponding attack altogether or apply a particular alert filter to that attack that will disable alerting for a particular IP address or range of IP addresses. In almost all cases, it is a best practice to implement the latter. For instance, an SMS server may be generating the alert **Netbios: Copy Executable file attempt** during the legitimate transfer of login scripts. Rather than disable the alert altogether, and cancel the possibility of finding a real attack of this nature, we recommend that you create an alert filter for the SMS server and applied to this attack.

Every alert filter created is globally stored so that the filter can be applied to any Exploit or Reconnaissance attack.

It is also a best practice to document all your tuning activities. The Report Generator can be used to assist the documentation process. The IDS Policy report will deliver reports that list Alert Filters that have been applied and attacks that have been otherwise customized.



See Chapter 11 of the *Manager Administrator's Guide* for more details on alert filters.

DoS

It is a best practice to let the sensors learn the profiles of the particular segments they are monitoring before tuning DoS attacks. This is Learning Mode operation. The learning process takes two days. During this period it is not unusual to see DoS alerts associated with normal traffic flows (e.g DoS SYN flood alerts reported outbound on a firewall interface to the Internet). After a profile has been learned, the particulars of the profile (number of SYNS, ACKS, etc.) can be viewed per sensor. DoS detection can also be implemented using the Threshold Mode. This involves setting thresholds manually for the type of segment characteristics that are learned in Learning Mode. Implementing this mode successfully is critically dependent on detailed knowledge of the segments the particular sensors are monitoring.

It is a best practice to have the sensor re-learn the profile when there is a network change (i.e., you move the sensor from a lab or staging environment to a production environment) or a configuration change (i.e., you change the CIDR block of a sub-interface) that causes a significant sudden traffic change on an interface. If the sensor does not re-learn the new environment, it may issue false alarms or fail to detect actual attacks during a time period when it is adapting to the new network traffic conditions. There is no need to re-learn a profile when network traffic increases or decreases naturally over time (e.g., an eCommerce site that is getting more and more customers; thus its Web traffic increases in parallel), since the sensor can automatically adapt to it.

See Chapter 9 of the *Manager Administrator's Guide* for more details about re-learn profile.

Sensor actions

There are multiple sensor actions that are available for configuration per attack. These include:

- **Dropping Further Packets:** Only works in in-line mode. Will drop a detected attack packet and all subsequent packets in the same flow.
- **Firewall Action:** Sensor will communicate with a designated firewall to dynamically configure ACL's. This feature works with a limited number of firewalls and is rarely used in the field. Future software releases will add the implementation of ACL's to sensor capabilities



See Chapter 7 of the *Manager Administrator's Guide* for more details on sensor actions.

Response management

When an IntruShield sensor detects activity to be in violation of a configured policy, a preset response from the sensor is integral to the protection or prevention process. Proper configuration of responses is crucial to maintaining effective protection. Critical attacks like buffer overflows and DoS attacks require responses in real time, while scans and probes can be logged and researched to determine compromise potential and the source of the attack.

Developing a system of actions, alerts, and logs based on specific attacks or attack parameters (such as severity) is recommended for effective network security. For example, since IntruShield can be customized to protect any zone in a network, knowing what needs to be protected can help to determine the response type. If monitoring outside of the firewall in in-line mode, preventing DoS attacks and attacks against the firewall is crucial. Most other suspicious traffic intended for the internal network, including scans and low-impact well-known exploits, are best logged and analyzed as the impact is not immediate and a better understanding of the potential attack purpose can be determined. Thus, if you are monitoring outside of a firewall in in-line mode, it is important to not set the policies and responses so fine that they disrupt the flow of traffic and slow down the system; rather, prevent the crippling traffic from disrupting your network.

Remember that response actions are decoupled from alerting. Pay particular attention to this with the Recommended For Blocking (RFB) category of attacks, lest you enable blocking for an attack, but disable alerting, causing the attack to be blocked without your knowledge. (Unless that is your goal.)

Creating rule sets

Proper creation of rule sets is essential to eliminating false positives and ensure maximum protection on your network. These best practices can assist when creating rules sets in the IntruShield Manager.

Default Inline IPS

A rule set is configured based on attack category, operating system, protocol, application, severity, and benign trigger probability options. Each rule in a set is either an include rule or an exclude rule. An include rule (which should always start a rule set) is a set of parameters that encompass a broad range of well-known attacks for detection. An exclude rule removes elements from the include rule in order to focus the policy's rule set.

There are two best practice methods employed for creating rule sets.

- **General-to-specific rule creation.** The first method is general-to-specific. Start with an include rule that covers a broad range of OSs, applications, protocols. After this, create one or more exclude rules to strip away specific OSs, protocols, et cetera, thus focusing the rule set on the environment where it will be enforced. For example, start with an include rule for all Exploit category attacks. Follow this with multiple exclusion rules that strip away protocols, applications, severities, et cetera, that are rarely or never seen in a zone of your network.
- **Collaborative rule creation.** The second method is collaboration: Create multiple include rules within one rule set for each category, OS, et cetera, combination that needs to be detected. Each criterion must be matched in order for an alert to be triggered. For example, create the first rule in the set with the Exploit category, Unix as the OS, Sendmail as the application, and SMTP as the protocol. Next, create another include rule for Exploit, Windows 2000, WindMail, and so forth in the same manner. Each include rule added broadens the scope of the detection.



See Chapter 7 of the *Manager Administrator's Guide* for more details about rule set creation.

Port clustering on asymmetric networks

Port clustering, referred to as Interface Groups in the IntruShield Manager interface, enables multiple ports on a single sensor to be grouped together for effective traffic monitoring. It is a best practice to implement a port clustering configuration when dealing with asymmetrically routed networks. Asymmetric networks are common in load balancing and active/passive configurations, and a complete transmission may be received on one segment, but depart on another. Thus keeping state of asymmetric transmissions is essential for successfully monitoring the traffic. Interface groups normalize the impact of traffic flows split across multiple interfaces, thus maintaining state to avoid information loss.

Once configured, an interface group appears in the System Configuration tool's Resource Tree as a single interface node (icon) under the sensor where the ports are located. All of the ports that make up the interface are configured as one logical entity, keeping the configuration consistent.



See Chapter 4 of the *Manager Administrator's Guide* for more details on port clustering.

Maintenance, backup, and database tuning

- Perform regular manual backups of your database using the Backup feature in the Manager software. Your configuration tables are saved by default once a week on Saturday. See [Backup on page 15](#), for more information on backup best practices.
- Database backups are cumulative and the size of a backup file can become quite large. Perform regular file maintenance to prevent disk space issues.
- Archive your alerts and packet logs regularly, using the Alert and Packet Log Archival feature. McAfee recommends that you archive your alert data monthly, and that you discard alert and packet log information from your database every 90 days to manage your database size. Note that there is currently a 4GB size limitation for a single archive file.



A database left untuned can, over time, lead to data corruption.

- Online database tuning operation causes the creation of temporary alert and packet log tables; if you are using an agent that queries the database, your agent may attempt to interact with these tables during tuning. There is a remote chance that during the transition to the temporary tables, the SQL query will result in an error. If a SQL query error occurs, simply retry the query. Further information on the impact of online database tuning of the LSM database will be sent to the third-party vendors that are directly accessing this database. If you have any specific questions, contact Technical Support. Also note that there is no change in database SQL query behavior if online database tuning is disabled.
- McAfee recommends that you make a regular practice of defragmenting the disk of the Manager server, as disk fragmentation can lead to database inefficiency.
- When scheduling certain Manager actions (backups, file maintenance, archivals, database tuning), set a time for each that is unique and is a minimum of an hour after/before other scheduled actions. Do not run scheduled actions concurrently.



See Chapter 6 and Appendix of the *Manager Administrator's Guide* for more information on tuning your database.

- McAfee recommends that you tune your database at regular intervals using the online tuning tools available beginning with release 2.1.1.
- Refer to the *Oracle9i Deployment Guide document* (included on the product CD) for information on Oracle hardware requirements, Oracle database setup, IntruShield Manager installation for use with a remote Oracle database, and for information on proper database tuning.
- If you are using an Oracle database, McAfee strongly recommends that you employ an Oracle DBA to maintain your database.

Alerts and disk space maintenance

Disk space maintenance is an important task that must be completed to ensure efficient running of the Manager. In order to develop best practices for database maintenance it is important to understand the lifecycle of an alert.

Alert states

Alerts exist in one of three states: unacknowledged/acknowledged, and marked for deletion. When an alert is raised, it appears in the Manager in an unacknowledged state. Unacknowledged means that you have not officially recognized its presence by marking it acknowledged. An alert remains in an unacknowledged state until you either acknowledge or delete it. Alerts are backed up to the database and archived in order of occurrence. Deleted alerts are removed from the database.

Unacknowledged alerts display in the Unacknowledged Alert Summary section of the Network Console and the Real-time view in the Alert Viewer. Acknowledging alerts dismisses them from these views. Acknowledged alerts shown only in the Historical view in the Alert Viewer and in reports.

Deleting an alert both acknowledges it and marks it for deletion. The alert is not actually deleted until a scheduled File Maintenance takes place. At that time, Intrushield deletes those alerts marked for deletion and those alerts meeting the deletion criteria specified in the scheduler-older than 30 days, for example-whether or not they've been manually marked for deletion.

To put an acknowledged alert back into an unacknowledged state or un-delete an alert, you can use the Historical view in Alert Viewer to show all alerts from the time period in which the acknowledged/deleted alert took place. You can then locate the alert and unacknowledge or un-delete it. This alert will not display in the Real-time Alert Viewer until you have closed and re-opened the Alert Viewer.

Databases can be substantial, possibly containing all Alert and Packet logs, any incident reports that have been generated and audit and fault logs. Maintenance of this data can be accomplished automatically using the File Maintenance scheduler. If you have a large amount of data and wish to do your tuning offline, it is a best practice to use the **purge.bat** (forces the deletion of old alerts) and **dbtuning.bat** (force the tuning of the database) scripts. To do this you must stop the Manager and run the scripts. To tune while the Manager is running, use the online tools, but ensure that you do not have another process running concurrently during the tuning.

Dbtuning.bat

The dbtuning.bat utility does the following:

- Defragments tables where rows/columns are split or have been deleted
- Re-sorts indexes
- Updates index statistics
- Computes query optimizer statistics
- Checks and repairs tables

Purge.bat

The purge.bat enables on-demand deletion of alerts and packet log data from your database. Alerts and packet logs can be deleted that are older than a specified number of days, or if they have been marked for deletion via the Alert Viewer tool.

Purge.bat also offers to automatically start dbtuning.bat immediately after the purge is completed.

Using the File Maintenance scheduler

A best-practice suggestion is to wait for 97 days of data and then on a recurring 7-day period run the purge.bat and dbtuning.bat—this will delete alerts already marked for deletion (from the Alert Viewer) as well as alerts older than 90 days. Scripts have to be run off-line (i.e., Manager service stopped) to release the lock from the database.

If automatic File Maintenance is used to delete alert and packet log data it is recommended that a large value—such as 90, as in 90 days—is entered in the “Scheduled Deletion” column for the Alert & Packet Log Data option. This allows for long-term analysis of alerts and logs without overburdening your database with millions of alerts, which may affect long-term and overall database performance. By setting the value to 90 days, all alerts and packet logs older than 90 days are deleted at the weekly maintenance scheduler time.

Apart from the database data, Intrushield Manager creates a group of administration files that must be maintained regularly. These include Diagnostic files, DoS files (profiles) and Data Mining files (for Trend Reporting) among others. It is a best practice to schedule the deletion of the oldest of these files on an on-going basis. This can be accomplished using the Maintenance scheduler.



See Chapter 6 and Appendix B of the *Manager Administrator's Guide* for more information on file maintenance.

Backup

McAfee recommends the following approach to backing up IntruShield data and configurations:

- Back up Manager data either within the Manager server (Intrushield\Backups folder) or preferably on external media.
- Back up all information, including configurations, alerts, and audits.
- Implement a schedule for backups using the Backup scheduler. Backing up config tables weekly is recommended. (Be sure to schedule this at a time when other processes will not be running concurrently.)
- As the 'All Tables' and 'Audit and Alert Tables' options can be rather large in size (depending upon the amount of alert data in the database) these types of backups should be saved off the Manager server.
- Saving the 'All Tables' settings monthly is strongly recommended.
- Protect backups from tampering by creating a digital fingerprint of the file using a hash function such as MD5 or SHA-1.

- Test restoration of backups periodically to ensure that a backup was successful and valid. The best way to do this is to perform a “test” restore of the backup on a secondary, non-production Manager.
- The 'Config Tables' option backs up only tabled information relating to configured tasks. This option is enabled by default to occur every Saturday night. This is set within the Backup Scheduler action.
- Save actual configs of sensors (not just the config tables) using the Export option under the Sensor_Name tab. This creates a xml file (no attempt to read this file should be made) that can be imported to any sensor of the same type in the future. Save actual sensor configs weekly.



See Chapter 6 and Appendix B of the *Manager Administrator's Guide* for more information on file maintenance.

Access Control Lists (ACL)

- When working with ACLs, note that you cannot set explicit ACL permit rules for protocols that negotiate ports dynamically, with the exception of FTP, TFTP, and RPC services. Protocols such as H.323 and Netmeeting, which negotiate the data channel separately from the control channel, or negotiate ports that do not follow a standard, are not supported. However, you can explicitly deny these protocol instances by denying the fixed control port. However, you can configure ACLs to explicitly deny these protocol instances by denying the fixed control port.
- For RPC services, you can configure explicit permit and deny rules for RPC as a whole, but not its constituents, such as statd and mountd.
- Protocols or services, such as instant messaging and peer-to-peer communication, that use dynamic ports, are not supported.
- An alternative option for denying protocols that use dynamic ports is to configure IDS policies to drop the attacks that are detected in such transmissions. IntruShield detects use of and attacks in such programs as Yahoo Messenger, KaZaA, IRC, and so on.
- There is a limit on the number of ACL rules that can be supported by a sensor.

Sensor	ACL rule limit
I-4010	1000
I-4000	1000
I-3000	1000
I-2700	400
I-2600	400
I-1400	100
I-1200	50



See Chapter 9 of the *Manager Administrator's Guide* for more details on ACLs.

Working on performance issues

Most performance issues are related to switch port configuration, duplex mismatches, link up/down situations, and data link errors.

Sniffer trace

A Sniffer details packet transfer, and thus a Sniffer trace analysis can help pinpoint switch and IntruShield performance or connectivity issues when the issues persist after you have exhausted the other suggestions in this document. Sniffer trace analysis reveals every packet on the wire and pinpoints the exact problem.

Note that it may be important to obtain several Sniffer traces from different ports on different switches, and that it is useful to monitor ("span") ports rather than spanning VLANs when troubleshooting switch connectivity issues.

Understanding data link errors

Many performance issues may be related to data link errors. Excessive errors usually indicate a problem. For more on this topic, see also [Ensuring connectivity between the sensor and other network devices on page 5](#).

Half-duplex setting

When operating with a duplex setting of half-duplex, some data link errors such as FCS, alignment, runts, and collisions are normal. Generally, a one percent ratio of errors to total traffic is acceptable for half-duplex connections. If the ratio of errors to input packets is greater than two or three percent, performance degradation may be noticeable.

In half-duplex environments, it is possible for both the switch and the connected device to sense the wire and transmit at exactly the same time, resulting in a collision. Collisions can cause runts, FCS, and alignment errors, which are caused when the frame is not completely copied to the wire, resulting in fragmented frames.

Full-duplex setting

When operating at full-duplex, FCS, cyclic redundancy checks (CRC), alignment errors, and runt counters should be minimal. If the link is operating at full-duplex, the collision counter is not active. If the FCS, CRC, alignment, or runt counters are incrementing, check for a duplex mismatch. Duplex mismatch is a situation in which the switch is operating at full-duplex and the connected device is operating at half-duplex, or vice versa. The result of a duplex mismatch is extremely slow performance, intermittent connectivity, and loss of connection. Other possible causes of data link errors at full-duplex are bad cables, a faulty switch port, or software or hardware issues.

SSL best practices

Note that there is a performance impact when using the SSL detection feature. The following are the SSL throughput measurements and test methodologies for the 2.1 release.

SSL only traffic - throughput

- Session resumption for 4 out of 5 TCP connections
- 5 HTTP 1.1 get page requests per TCP connection with a 10K response each
- 1024-bit RSA
- 128-bit ARC4

	I-2600	I-2700	I-3000	I-4000	I-4010
Max. SSL Connections / Sec.	200	300	400	800	800
Throughput	96 Mbps	150 Mbps	200 Mbps	400 Mbps	400 Mbps

SSL only traffic - throughput

- Session resumption for 4 out of 5 TCP connections
- 5 HTTP 1.1 get page requests per TCP connection with a 5K response each
- 1024-bit RSA
- 128-bit ARC4

	I-2600	I-2700	I-3000	I-4000	I-4010
Max. SSL Connections / Sec.	200	325	600	800	1200
Throughput	50 Mbps	85 Mbps	155 Mbps	200 Mbps	310 Mbps

SSL traffic mixed with HTTP 1.1 traffic

- Session resumption for 4 out of 5 TCP connections
- 5 HTTP 1.1 get page requests per TCP connection with a 5K response each
- 1024-bit RSA
- 128-bit ARC4

	I-2600	
Max. SSL Connections / Sec.	100	200
SSL Throughput	25 Mbps	50 Mbps
HTTP 1.1 Throughput	475 Mbps	350 Mbps
Total Throughput	500 Mbps	400 Mbps

I-2600		
Max. SSL Connections / Sec.	100	200
SSL Throughput	25 Mbps	50 Mbps
HTTP 1.1 Throughput	475 Mbps	350 Mbps
Total Throughput	500 Mbps	400 Mbps

I-3000		
Max. SSL Connections / Sec.	200	400
SSL Throughput	50 Mbps	105 Mbps
HTTP 1.1 Throughput	860 Mbps	475 Mbps
Total Throughput	910 Mbps	580 Mbps

I-4000		
Max. SSL Connections / Sec.	400	800
SSL Throughput	100 Mbps	200 Mbps
HTTP 1.1 Throughput	1550 Mbps	780 Mbps
Total Throughput	1650 Mbps	980 Mbps

I-4010		
Max. SSL Connections / Sec.	400	800
SSL Throughput	100 Mbps	200 Mbps
HTTP 1.1 Throughput	1740 Mbps	860 Mbps
Total Throughput	1840 Mbps	1060 Mbps

Supported Web servers

SSL decryption is supported for the following web servers:

- Microsoft Internet Information Server (IIS)
- Apache

Supported cipher suites

The following SSL cipher suites (as named in their respective RFCs) are supported:

- SSLv2 cipher suites
- SSL_CK_RC4_128_WITH_MD5
- SSL_CK_RC4_128_EXPORT40_WITH_MD5
- SSL_CK_DES_64_CBC_WITH_MD5
- SSL_CK_DES_192_EDE3_CBC_WITH_MD5

SSLv3/TLS cipher suites

- DES_CBC3_SHA
- TLS_NULL_WITH_NULL_NULL

- TLS_RSA_WITH_NULL_MD5
- TLS_RSA_WITH_NULL_SHA
- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_DES_CBC_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA

Unsupported SSL functionality

The following SSL functionalities are not supported:

- iPlanet Web servers
- Diffie-Hellman ciphers (McAfee recommends that you disable acceptance of Diffie-Hellman requests on the SSL Web server to ensure that IntruShield is able to decrypt the traffic)
- Compression in the SSL records (a negotiable option in SSLv3 and TLS)
- PCT (Microsoft's extension to SSLv2)

Sensor performance with HTTP Response processing

HTTP response processing is disabled by default. You can enable it for each traffic direction on an interface pair. To minimize the potential performance impact on the sensor, we recommend that you enable HTTP response processing on the minimum number of ports and in only the required directions to achieve your protection goals.

Some examples of HTTP response processing deployment:

- You want to protect a bunch of clients on your internal network – enable HTTP response processing for inbound traffic only.
- You are serving Web content and do not wish to serve attacks embedded in HTTP response traffic – enable HTTP response processing for outbound traffic only.
- You want to protect both clients internally and content you're serving. Enable HTTP response processing in both directions.

HTTP Response processing enabled for both inbound and outbound traffic

The following table shows sensor performance with HTTP response processing enabled in both directions:

- 5 HTTP 1.1 get page requests per TCP connection with a 5K response each sent in both directions

	I-1200	I-1400	I-2600	I-2700	I-3000	I-4000	I-4010
Max. aggregate HTTP traffic (across entire sensor)	72Mbps	129Mbps	200Mbps	210Mbps	430Mbps	740Mbps	860Mbps

HTTP Response processing enabled in one direction only (inbound OR outbound)

The following table shows sensor performance with HTTP response processing enabled for a single direction:

- 5 HTTP 1.1 get page requests per TCP connection with a 5K response each sent in both directions

	I-1200	I-1400	I-2600	I-2700	I-3000	I-4000	I-4010
Max. aggregate HTTP traffic (across entire sensor)	100Mbps	161Mbps	320Mbps	390Mbps	776Mbps	1150Mbps	1350Mbps